

CLAIMS

1. A method of accessing a unit of data stored in a storage environment that includes a plurality of storage clusters, the method comprising acts of:
receiving a request from a host computer to locate the unit of data previously stored in the storage environment; and
in response to receipt of the request, determining on which one of the plurality of storage clusters the unit of data is stored.
2. The method of claim 1, wherein the host computer executes an application program that stores data to and retrieves data from the storage environment, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the act of receiving is performed by the application programming interface.
3. The method of claim 1, wherein the storage environment is coupled to the host computer by at least one communication link, wherein the host computer, the storage environment, and the at least one communication link form a computer system, wherein the computer system includes an appliance that monitors access requests from the host computer, and wherein the act of receiving the request further comprises an act of receiving, at the appliance, the request from the application.
4. The method of claim 3, wherein the act of determining is performed by the appliance.
5. The method of claim 1, wherein the host computer executes an application program that stores data on the storage environment, and wherein the act of receiving the request further comprises an act of receiving the request directly from the application program.
6. The method of claim 4, wherein the act of receiving the request further comprises an act of receiving the request at at least one of the plurality of storage clusters.

7. The method of claim 6, wherein the at least one of the plurality of storage clusters includes at least one access node that receives and processes access requests, and wherein the act of receiving the request from the application program at the at least one of the plurality of storage clusters further comprises an act of receiving the request at the at least one access node.

8. The method of claim 1, wherein the act of determining comprises an act of performing a search for the unit of data on the plurality of storage clusters.

9. The method of claim 8, wherein the act of performing a search for the unit of data further comprises an act of performing the search serially through the plurality of storage clusters until the unit of data is found.

10. The method of claim 8, wherein the act of performing a search for the unit of data further comprises an act of performing the search on each of the plurality of storage clusters in parallel.

11. The method of claim 1, wherein the act of determining is performed by at least one of the plurality of storage clusters.

12. The method of claim 1, wherein the act of determining comprises locating the unit of data on at least one of the plurality of storage clusters without performing a search.

13. The method of claim 1, wherein the storage environment is a content addressable storage environment in which the unit of data stored in the storage environment is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data, and wherein the act of determining further comprises determining on which of the plurality of storage clusters the unit of data is stored based on the content address of the unit of data.

14. The method of claim 13, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the act of determining comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the time information of the content address of the unit of data.

15. The method of claim 14, wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash value of the time information of the content address of the unit of data.

16. The method of claim 13, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the GUID.

17. The method of claim 16, wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash of the GUID.

18. The method of claim 13, wherein the act of determining further comprises acts of:
accessing information that specifies an algorithm that was used to select on which of the plurality of storage clusters the unit of data was stored, based on the content address of the unit of data; and
applying the algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters the unit of data is stored.

19. The method of claim 18, wherein the information specifies a plurality of algorithms used by the storage environment and during which period of time each of the plurality of algorithms was used to store units of data.

20. The method of claim 19, wherein the information further specifies, for each one of the plurality of algorithms, at least one storage cluster that was in the storage environment during the period of time when the one of the plurality of algorithms was in effect.

21. At least one computer readable medium encoded with instructions that, when executed on a computer system, perform a method of accessing a unit of data stored in a storage environment that includes a plurality of storage clusters, wherein the computer system comprises the storage environment and a host computer that stores data on the storage environment, and wherein the method comprises acts of:

receiving a request from the host computer to locate the unit of data previously stored in the storage environment; and

in response to receipt of the request, determining on which one of the plurality of storage clusters the unit of data is stored.

22. The at least one computer readable medium of claim 21, wherein the host computer executes an application program that stores data to and retrieves data from the storage environment, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the act of receiving is performed by the application programming interface.

23. The at least one computer readable medium of claim 21, wherein the computer system further comprises at least one communication link that couples the storage environment to the host computer and an appliance that monitors access requests from the host computer to the storage environment, and wherein the act of receiving the request further comprises an act of receiving, at the appliance, the request from the application.

24. The at least on computer readable medium of claim 23, wherein the act of detecting is performed by the appliance.

25. The at least one computer readable medium of claim 21, wherein the host computer executes an application program that stores data on the storage environment, and wherein the act of receiving the request further comprises an act of receiving the request directly from the application program.

26. The at least one computer readable medium of claim 23, wherein the act of receiving the request further comprises an act of receiving the request at at least one of the plurality of storage clusters.

27. The at least one computer readable medium of claim 26 wherein the at least one of the plurality of storage clusters includes at least one access node that receives and processes access requests, and wherein the act of receiving the request from the application program at the at least one of the plurality of storage clusters further comprises an act of receiving the request at the at least one access node.

28. The at least one computer readable medium of claim 21, wherein the act of determining comprises an act of performing a search for the unit of data on the plurality of storage clusters.

29. The at least one computer readable medium of claim 28, wherein the act of performing a search for the unit of data further comprises an act of performing the search serially through the plurality of storage clusters until the unit of data is found.

30. The at least one computer readable medium of claim 28, wherein the act of performing a search for the unit of data further comprises an act of performing the search on each of the plurality of storage clusters in parallel.

31. The at least one computer readable medium of claim 21, wherein the act of determining is performed by at least one of the plurality of storage clusters.

32. The at least one computer readable medium of claim 21, wherein the act of determining comprises locating the unit of data on at least one of the plurality of storage clusters without performing a search.

33. The at least one computer readable medium of claim 21, wherein the storage environment is a content addressable storage environment in which the unit of data stored in the storage environment is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data, and wherein the act of determining further comprises determining on which of the plurality of storage clusters the unit of data is stored based on the content address of the unit of data.

34. The at least one computer readable medium of claim 33, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the act of determining comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the time information of the content address of the unit of data.

35. The at least one computer readable medium of claim 34, wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash value of the time information of the content address of the unit of data.

36. The at least one computer readable medium of claim 33, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the GUID.

37. The at least one computer readable medium of claim 36, wherein the act of determining further comprises an act of determining on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash of the GUID.

38. The at least one computer readable medium of claim 33, wherein the act of determining further comprises acts of:
- accessing information that specifies an algorithm that was used to select on which of the plurality of storage clusters the unit of data was stored, based on the content address of the unit of data; and
 - applying the algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters the unit of data is stored.
39. The at least one computer readable medium of claim 38, wherein the information specifies a plurality of algorithms used by the storage environment and during which period of time each of the plurality of algorithms was used to store units of data.
40. The at least one computer readable medium of claim 39, wherein the information further specifies, for each one of the plurality of algorithms, at least one storage cluster that was in the storage environment during the period of time when the one of the plurality of algorithms was in effect.
41. An apparatus for storing data, the apparatus comprising:
- an input that receives a request from a host computer to locate a unit of data on at least one of a plurality of storage clusters in a storage environment, and
 - at least one controller, coupled to the input, that:
 - receives the request from the input; and
 - in response to receipt of the request, determines on which of the plurality of storage clusters the unit of data is stored.
42. The apparatus of claim 41 in combination with the storage environment, the host computer that accesses data stored in the storage environment; and a communication link that couples the host computer to the storage environment to form a computer system.
43. The apparatus of claim 42, wherein the at least one controller is disposed in the host computer.

44. The apparatus of claim 42, wherein the at least one controller is disposed in the storage environment.
45. The apparatus of claim 42 wherein the at least one controller is disposed in between the storage environment and the host computer in an appliance that monitors access requests from the host computer to the storage environment.
46. The apparatus of claim 41, wherein the host computer executes an application program that stores data to and retrieves data from the storage environment, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the at least one controller receives the request at the application programming interface.
47. The apparatus of claim 41, wherein the host computer executes an application program that stores data in the storage environment, and wherein the at least one controller receives the request directly from the application program.
48. The apparatus of claim 44, wherein the apparatus is disposed in at least one of the plurality of storage clusters.
49. The apparatus of claim 48, wherein the at least one of the plurality of storage clusters includes at least one access node that receives and processes access requests, and wherein the at least one controller receives the request at the at least one access node.
50. The apparatus of claim 41, wherein the at least one controller performs a search for the unit of data on the plurality of storage clusters.
51. The apparatus of claim 50, wherein the at least one controller performs the search serially through the plurality of storage clusters until the unit of data is found.
52. The apparatus of claim 50, wherein the at least one controller performs the search on each of the plurality of storage clusters in parallel.

53. The apparatus of claim 41, wherein the at least one controller performs the search at at least one of the plurality of storage clusters.

54. The apparatus of claim 41, wherein the at least one controller locates the unit of data on at least one of the plurality of storage clusters without performing a search.

55. The apparatus of claim 41, wherein the storage environment is a content addressable storage environment in which the unit of data stored in the storage environment is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data, and wherein the at least one controller determines on which of the plurality of storage clusters the unit of data is stored based on the content address of the unit of data.

56. The apparatus of claim 55, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the at least one controller determines on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the time information of the content address of the unit of data.

57. The apparatus of claim 56, wherein the at least one controller determines on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash value of the time information of the content address of the unit of data.

58. The apparatus of claim 57, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the at least one controller determines on which of the plurality of storage clusters the unit of data is stored based, at least in part, on the GUID.

59. The apparatus of claim 58, wherein the at least one controller determines on which of the plurality of storage clusters the unit of data is stored based, at least in part, on a hash of the GUID.

60. The apparatus of claim 55, wherein the at least one controller:
accesses information that specifies an algorithm that was used to select on which
of the plurality of storage clusters the unit of data was stored, based on the content
address of the unit of data; and

applies the algorithm to the content address of the unit of data to determine on
which of the plurality of storage clusters the unit of data is stored.

61. The apparatus of claim 60, wherein the information specifies a plurality of
algorithms used by the storage environment and during which period of time each of the
plurality of algorithms was used to store units of data.

62. The apparatus of claim 61, wherein the information further specifies, for each one
of the plurality of algorithms, at least one storage cluster that was in the storage
environment during the period of time when the one of the plurality of algorithms was in
effect.

63. A method of accessing a unit of data stored in a storage environment that includes
a plurality of storage clusters, the storage environment storing data for a host computer,
the method comprising acts of:

receiving a request from an application program executing on the host computer
to store a unit of data; and

in response to receipt of the request, selecting one of the plurality of storage
clusters to store the unit of data in a manner that is transparent to the application program
so that the application program is not made aware that the selected one of the plurality of
storage clusters stores the unit of data.

64. The method of claim 63, wherein the host further executes an application
programming interface that interfaces the application program to the storage
environment, and wherein the act of receiving is performed by the application
programming interface.

65. The method of claim 63, wherein the storage environment is coupled to the host computer by at least one communication link, wherein the host computer, the storage environment, and the at least one communication link form a computer system, wherein the computer system includes an appliance that monitors access requests from the host computer, and wherein the act of receiving the request further comprises an act of receiving, at the appliance, the request from the application.

66. The method of claim 63, wherein the act of receiving the request further comprises an act of receiving the request directly from the application program.

67. The method of claim 66, wherein the act of receiving the request from the application program further comprises an act of receiving the request at at least one of the plurality of storage clusters.

68. The method of claim 63, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data using a round-robin technique.

69. The method of claim 63, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on a load of at least one of the plurality of storage clusters.

70. The method of claim 63, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on an available storage capacity of each of the plurality of storage clusters.

71. The method of claim 63, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on a size of the unit of data.

72. The method of claim 63, wherein the storage environment is a content addressable storage environment in which the unit of data stored in the storage

environment is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data, and wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on the content address of the unit of data.

73. The method of claim 72, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the act of selecting comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on the time information of the content address of the unit of data.

74. The method of claim 73, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash value of the time information of the content address of the unit of data.

75. The method of claim 72, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on the GUID.

76. The method of claim 75, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash of the GUID.

77. The method of claim 72, wherein the act of selecting further comprises acts of:
applying an algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters to store the unit of data;
storing the algorithm in a record that indicates a time frame in which the algorithm was in use.

78. The method of claim 77, wherein the record specifies a plurality of algorithms used by the storage environment and at what period of time each of the plurality of algorithms was used to store units of data.

79. The method of claim 63, further comprising an act of:
storing the unit of data on the selected one of the plurality of clusters.

80. At least one computer readable medium encoded with instructions that, when executed on a computer system, perform a method of accessing a unit of data stored in a storage environment that includes a plurality of storage clusters, wherein the computer system comprises the storage environment and a host computer that stores data on the storage environment, and wherein the method comprises acts of:

receiving a request from an application program executing on the host computer to store a unit of data; and

in response to receipt of the request, selecting one of the plurality of storage clusters as to store the unit of data in a manner that is transparent to the application program so that the application program is not made aware that the selected one of the plurality of storage clusters stores the unit of data.

81. The at least one computer readable medium of claim 80, wherein the host further executes an application programming interface that interfaces the application program to the storage environment, and wherein the act of receiving is performed by the application programming interface.

82. The at least one computer readable medium of claim 80, wherein the computer system further comprises at least one communication link that couples the storage environment to the host computer to the storage environment and an appliance that monitors access requests from the host computer, and wherein the act of receiving the request further comprises an act of receiving, at the appliance, the request from the application.

83. The at least one computer readable medium of claim 80, wherein the act of receiving the request further comprises an act of receiving the request directly from the application program.

84. The at least one computer readable medium of claim 83, wherein the act of receiving the request from the application program further comprises an act of receiving the request at at least one of the plurality of storage clusters.

85. The at least one computer readable medium of claim 80, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data using a round-robin technique.

86. The at least one computer readable medium of claim 80, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on an available storage capacity of each of the plurality of storage clusters.

87. The at least one computer readable medium of claim 80, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on a load of at least one of the plurality of storage clusters.

88. The at least one computer readable medium of claim 80, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on a size of the unit of data.

89. The at least one computer readable medium of claim 80, wherein the storage environment is a content addressable storage environment in which the unit of data stored in the storage environment is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data, and wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based on the content address of the unit of data.

90. The at least one computer readable medium of claim 89, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and the act of selecting comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on the time information of the content address of the unit of data.

91. The at least one computer readable medium of claim 90, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash value of the time information of the content address of the unit of data.

92. The at least one computer readable medium of claim 89, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on the GUID.

93. The at least one computer readable medium of claim 91, wherein the act of selecting further comprises an act of selecting one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash of the GUID.

94. The at least one computer readable medium of claim 89, wherein the act of selecting further comprises acts of:

applying an algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters to store the unit of data;

storing the algorithm in a record that indicates a time frame in which the algorithm was in use.

95. The at least one computer readable medium of claim 94, wherein the record specifies a plurality of algorithms used by the storage environment and during which period of time each of the plurality of algorithms was used to store units of data.

96. The at least one computer readable medium of claim 80, wherein the method further comprises an act of:
storing the unit of data on the selected one of the plurality of clusters.
97. An apparatus for storing data, the apparatus comprising:
an input that receives a request from an application program executing on a host computer to store a unit of data on at least one of a plurality of storage clusters in a storage environment, and
at least one controller, coupled to the input, that:
receives the request from the input; and
in response to receipt of the request, selects one of the plurality of storage clusters to store the unit of data in a manner that is transparent to the application program so that the application program is not made aware that the selected one of the plurality of storage clusters stores the unit of data.
98. The apparatus of claim 97 in combination with the storage environment, the host computer that accesses data stored in the storage environment; and a communication link that couples the host computer to the storage environment to form a computer system.
99. The apparatus of claim 98, wherein the at least one controller is disposed in the host computer.
100. The apparatus of claim 98, wherein the at least one controller is disposed in the storage environment.
101. The apparatus of claim 98 wherein the at least one controller is disposed in between the storage environment and the host computer in an appliance that monitors access requests from the host computer to the storage environment.
102. The apparatus of claim 99, wherein the host further executes an application programming interface that interfaces the application program to the storage

environment, and wherein the at least one controller receives the request at the application programming interface.

103. The apparatus of claim 97, wherein the at least one controller receives the request directly from the application program through the input.

104. The apparatus of claim 100, wherein the at least one controller receives the request at at least one of the plurality of storage clusters.

105. The apparatus of claim 97, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data using a round-robin technique.

106. The apparatus of claim 97, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on an available storage capacity of each of the plurality of storage clusters.

107. The apparatus of claim 97, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based on a load of at least one of the plurality of storage clusters.

108. The apparatus of claim 97, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based on a size of the unit of data.

109. The apparatus of claim 97, wherein the storage environment is a content addressable storage environment in which the unit of data stored in the storage environment is accessible by a content address that is based, at least in part, upon at least a portion of the content of the unit of data, and wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based on the content address of the unit of data.

110. The apparatus of claim 109, wherein the content address includes time information, based on when the unit of data was stored in the storage environment, and

the at least one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on the time information of the content address of the unit of data.

111. The apparatus of claim 110, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash value of the time information of the content address of the unit of data.

112. The apparatus of claim 109, wherein the content address includes a guaranteed unique identifier (GUID), and wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on the GUID.

113. The apparatus of claim 112, wherein the at least one controller selects one of the plurality of storage clusters to store the unit of data based, at least in part, on a hash of the GUID.

114. The apparatus of claim 109, wherein the at least one controller: applies an algorithm to the content address of the unit of data to determine on which of the plurality of storage clusters to store the unit of data; and stores the algorithm in a record that indicates a time frame in which the algorithm was in use.

115. The apparatus of claim 114, wherein the record specifies a plurality of algorithms used by the storage environment and during which period of time each of the plurality of algorithms was used to store units of data.

116. The apparatus of claim 97, wherein the at least one controller stores the unit of data on the selected one of the plurality of clusters.